arXiv:1704.02646v1 [math.ST] 9 Apr 2017

Posterior Asymptotic Normality for an Individual Coordinate in High-dimensional Linear Regression

Dana Yang*

Department of Statistics and Data Science, Yale University

Abstract

We consider the sparse high-dimensional linear regression model $Y = Xb + \epsilon$ where b is a sparse vector. For the Bayesian approach to this problem, many authors have considered the behavior of the posterior distribution when, in truth, $Y = X\beta + \epsilon$ for some given β. There have been numerous results about the rate at which the posterior distribution concentrates around β, but few results about the shape of that posterior distribution. We propose a prior distribution for b such that the marginal posterior distribution of an individual coordinate b_i is asymptotically normal centered around an asymptotically efficient estimator, under the truth. Such a result gives Bayesian credible intervals that match with the confidence intervals obtained from an asymptotically efficient estimator for b_i. We also discuss ways of obtaining such asymptotically efficient estimators on individual coordinates. We compare the two-step procedure proposed by Zhang and Zhang [9] and a one-step modified penalization method.

1 Introduction.

Consider the regression model

$$Y = Xb + \epsilon, \quad \epsilon \sim \mathcal{N}(0, I_n). \tag{1}$$

The design matrix X is of dimension $n \times p$. We are particularly interested in the case where $p > n$, for which b itself is not identifiable. In such a setting identifiability can be attained by adding a sparsity constraint on $|b|_0$, the number of nonzero b_i's. That is, the model consists of a family of probability measures $\{\mathbb{P}_b : b \in \mathbb{R}^p, |b|_0 \leq s^*\}$, and the observation Y is distributed $\mathcal{N}(Xb, I_n)$ under $\mathbb{P}_b$.

We are interested in the Bayesian inference on the vector b, when Y is actually distributed $N(X\beta, I_n)$ for some truth β. If p were fixed and X were

*xiaoqian.yang@yale.edu

full rank, classical theorems (the Bernstein-von Mises theorem, as in [8, page 141]) gives conditions under which the posterior distribution of b is asymptotically normal centered at the least squares estimator, with variance $(X^TX)^{-1}$ under $\mathbb{P}_\beta$.

The classical theorem fails when $p > n$. Although sparse priors have been proposed that give good posterior contraction rates [3] [5], posterior normality of b is only obtained under strong signal-to-noise ratio (SNR) conditions, such as the SNR conditions of Castillo el al. [3, Corollary 2], which forced the posterior to eventually have the same support as β. Effectively, their conditions reduce the problem to the classical, fixed dimensional case. However that is not the most interesting scenario. Without the SNR condition, Castillo et al. [3, Theorem 6] pointed out that under the sparse prior, the posterior distribution of b behaves like a mixture of Gaussians.

However, there is hope to obtain posterior normality results without the SNR condition if one considers the situation where only one component of b is of interest, say b_1, without loss of generality. All the other components are viewed as nuisance parameters. As shown by Zhang and Zhang [9] in a non-Bayesian setting, it is possible to construct estimators that are efficient in the classical sense that

$$\hat{\beta}_1 = \beta_1 + \frac{X_1^T \epsilon}{|X_1|^2} + o_p\left(\frac{1}{\sqrt{n}}\right). \tag{2}$$

We will use $o_p(\cdot)$ as a short hand for a stochastically small order term under $\mathbb{P}_\beta$ throughout this document. Here X_i denotes the i'th column of X, and the $o_p(\cdot)$ indicates that a term is of stochastically smaller order under $\mathbb{P}_\beta$. Later we also writeX_{-i} to denote the $n \times (p-1)$ matrix formed by all columns of X except for X_i. The $|\cdot|$ norm on a vector refers to the Euclidean norm.

Approximation (2) is useful when $|X_1|$ is of order $\sqrt{n}$, in which the expansion (2) implies weak convergence [6, page 171]:

$$|X_1|(\hat{\beta}_1^{(ZZ)} - \beta_1) \rightsquigarrow \mathcal{N}(0,1)$$

under $\mathbb{P}_\beta$ (Such behavior for $|X_1|$ is obtained with high probability when X is generated *i.i.d.* from the standard normal distribution). More precisely, Zhang and Zhang [9] proposed a two-step estimator that satisfies (2) under some regularity assumptions on X and no SNR conditions. They required the following behavior for X.

Assumption 1. *Let* $\gamma_i = X_1^T X_i/|X_1|^2, \lambda_n = \sqrt{\frac{\log p}{n}}$. *There exists a constant* $c_1 > 0$ *for which*

$$\max_{2\leq i\leq p} \gamma_i \leq c_1\lambda_n.$$

Assumption 2. *(REC*$(3s^*, c_2)$*) There exists constants* $c_2, c' > 0$ *for which*

$$\kappa(3s^*, c_2) = \min_{\substack{J\subset[p],\\ |J|\leq 3s^*}} \inf_{\substack{b\neq 0,\\ |b_{J^C}|_1\leq c_2|b_J|_1}} \frac{|Xb|}{\sqrt{n}|b_J|} > c' > 0. \tag{3}$$

Assumption 3. *The model dimension satisfies*

$$s^* \log p = o(\sqrt{n}).$$

Remark 1. *Assumption 2 is known as the restricted eigenvalue condition [1, page 1710] required for penalized regression estimators such as the LASSO estimator [7, page 1] and the Dantzig selector [2, page 1] to enjoy optimal* l^1 *and* l^2 *convergence rates.*

Theorem 1. *[9, Section 2.1,3.1] Under assumptions 1, 2 and 3, the estimator* $\hat{\beta}_1^{(ZZ)}$ *has expansion* (2).

The exact form of the estimator $\hat{\beta}_1^{(ZZ)}$ will be given in section 2.1.

The goal of this paper is to give a Bayesian analogue for Theorem 1, in the form of a prior distribution on b such that as $n, p \to \infty$, the posterior distribution of b_1 starts to resemble a normal distribution to centered around an estimator in the form of (2). Note that the sparse prior introduced by Castillo et al. [3] does not meet our goal since the marginal posterior distribution of b_1 under the sparse prior converges weakly to a mixture of normal distributions without consistent model selection.

Theorem 2. *Under assumptions 1, 2 and 3 and the constraint* $|X_1| = O(\sqrt{n})$, *there exists a prior on* b *for which the posterior distribution of* $|X_1|(b_1 - \hat{\beta}_1)$ *satisfies*

$$\left\| \mathcal{L}\left(|X_1|(b_1 - \hat{\beta}_1)|Y\right) - \mathcal{N}(0,1) \right\|_{BL} \to 0 \text{ in } \mathbb{P}_\beta, \tag{4}$$

where $\hat{\beta}_1$ *is an estimator of* β_1 *with expansion* (2).

The measure used here to quantify the discrepancy between probability measures is the bounded-Lipschitz metric [4, page 1]. The convergence of a sequence of distributions to a fixed distribution in bounded-Lipschitz metric is equivalent to weak convergence.

2 The prior and its background stories.

2.1 How does de-biasing work?

In sparse linear regression, penalized likelihood estimators such as the LASSO are often used and tend to give good global properties. One desirable property is the following bound on the l_1 loss.

$$\mathbb{P}_\beta \left\{ |\tilde{\beta} - \beta|_1 > Cs^* \lambda_n \right\} \to 0 \text{ as } n, p \to \infty \text{ for some } C > 0, \tag{5}$$

where λ_n is as defined in assumption 1. For example, Bickel et al. [1, Theorem 7.1] showed that under the REC condition (assumption 2) the LASSO estimator satisfies (5).

In general, penalized likelihood estimators introduce bias for the estimation of individual coordinates. To eliminate this bias, Zhang and Zhang [9] proposed a two-step procedure. First find a $\tilde{\beta}$, perhaps via a LASSO procedure that satisfies (5). Then define

$$\hat{\beta}_1^{(ZZ)} = \arg\min_{b_1 \in \mathbb{R}} |Y - X_{-1}\tilde{\beta}_{-1} - b_1 X_1|^2,$$

The idea behind this estimator is to penalize the magnitude of all coordinates except the one of interest. Under assumptions 1, 2 and 3, the one-step estimator $\hat{\beta}_1^{(ZZ)}$ is asymptotically unbiased with expansion (2). The same asymptotic behavior can be obtained in a single step, as in the next theorem. The idea of penalizing all coordinates but one to eliminate the bias is seen more clearly here.

Theorem 3. *Define*

$$\hat{\beta}^{(one-step)} = \arg\min_{b \in \mathbb{R}^p} \left(|Y - Xb|_2^2 + \eta_n \sum_{i \geq 2} |b_i| \right).$$

Choose η_n to be a large enough multiple of $n\lambda_n$. Under assumptions 1, 2 and 3, the one-step de-biasing estimator of β_1 achieves l_1 control (5) *and de-biasing simultaneously. The estimator for the first coordinate satisfies*

$$\hat{\beta}_1^{(one-step)} = \beta_1 + \frac{X_1^T \epsilon}{|X_1|^2} + o_p\left(\frac{1}{\sqrt{n}}\right).$$

Proof. In the proof of theorem 3 we will refer to the one step estimator as $\hat{\beta}$. We will first show that $\hat{\beta}$ satisfies (5). We know that when the penalty

involves all coordinates of b, then the bound on the l_1 norm is true [1, Theorem 7.1]. It turned out that leaving one term out the of penalty does not ruin that property.

As in the proof of [1, Theorem 7.1], we compare the evaluation of the penalized likelihood function at $\hat{\beta}$ and the truth β using the definition of $\hat{\beta}$.

$$|Y - X\hat{\beta}|_2^2 + \eta_n|\hat{\beta}_{-1}|_1 \leq |Y - X\beta|_2^2 + \eta_n|\beta_{-1}|_1.$$

Plug in $Y = X\beta + \epsilon$, the above is reduced to

$$|X(\hat{\beta} - \beta)|_2^2 \leq 2\sum_{i \leq n} \xi_i(\hat{\beta}_i - \beta_i) + \eta_n(|\beta_{-1}|_1 - |\hat{\beta}_{-1}|_1),$$

where $\xi_i = X_i^T\epsilon$. With high probability $|\max_{i \leq n} \xi_i| \leq R = C_2 n\lambda_n$, in which case we have

$$|X(\hat{\beta} - \beta)|_2^2 \leq 2R|\hat{\beta} - \beta|_1 + \eta_n(|\beta_{-1}|_1 - |\hat{\beta}_{-1}|_1). \tag{6}$$

From here we need to discuss two situations. First consider the case where 1 is in the S, the support of β. The expression above is bounded by

$$(2R + \eta_n)|(\hat{\beta} - \beta)_S|_1 + (2R - \eta_n)|(\hat{\beta} - \beta)_{S^C}|_1.$$

By choosing η_n to be a large enough multiple of $n\lambda_n$, we have

$$|X(\hat{\beta} - \beta)|_2^2 \leq c_1 n\lambda_n|(\hat{\beta} - \beta)_S|_1 - c_2 n\lambda_n|(\hat{\beta} - \beta)_{S^C}|_1.$$

Since the lefthand side is nonnegative, the above implies

$$|(\hat{\beta} - \beta)_{S^C}|_1 \leq \frac{c_1}{c_2}|(\hat{\beta} - \beta)_S|_1. \tag{7}$$

Therefore under assumption REC(c_0/c_1, κ), we can further bound the prediction loss by

$$\begin{aligned} &c_1\sqrt{s^* n \log p} \cdot |(\hat{\beta} - \beta)_S|_2 \\ \leq &\frac{c_1}{\kappa}\sqrt{s^* \log p} \cdot |X(\hat{\beta} - \beta)|_2. \end{aligned}$$

So far we have shown with high probability,

$$|X(\hat{\beta} - \beta)|_2 \leq \frac{c_1}{\kappa}\sqrt{s^* \log p}.$$

Under the REC assumption, we can go back to bound the $l1$ loss.

$$|(\hat{\beta}-\beta)_S|_1 \leq \sqrt{s^*}|(\hat{\beta}-\beta)_S|_2 \leq \frac{1}{\kappa}\sqrt{\frac{s^*}{n}}|X(\hat{\beta}-\beta)|_2 \leq \frac{c_1}{\kappa^2}s^*\lambda_n.$$

Therefore with (7) we have

$$|\hat{\beta}-\beta|_1 \leq \left(1+\frac{c_1}{c_2}\right)\frac{c_1}{\kappa^2}s^*\lambda_n.$$

The proof for the other case turned out to be messier. But the general idea remains the same. When $1 \in S^C$, we can bound the RHS of (6) by

$$(2R+\eta_n)\left|(\hat{\beta}-\beta)_{S\cup\{1\}}\right|_1 + (2R-\eta_n)\left|(\hat{\beta}-\beta)_{S^C\setminus\{1\}}\right|_1,$$

Choosing λ to be a large multiple of $\sqrt{n\log p}$ as in the $1 \in S$ case, we have

$$|X(\hat{\beta}-\beta)|_2^2 \leq c_1 n\lambda_n|(\hat{\beta}-\beta)_{S\cup\{1\}}|_1 - c_2 n\lambda_n|(\hat{\beta}-\beta)_{S^C\setminus\{1\}}|_1,$$

which implies

$$|(\hat{\beta}-\beta)_{S^C\setminus\{1\}}|_1 \leq \frac{c_1}{c_2}|(\hat{\beta}-\beta)_{S\cup\{1\}}|_1. \tag{8}$$

Again use assumption 2 to deduce the l_1 control (5).
Observe that the penalty term does not involve b_1.

$$\begin{aligned}\hat{\beta}_1 =& \arg\min_{b_1\in\mathbb{R}}|Y - X_{-1}\hat{\beta}_{-1} - b_1X_1|_2^2\\ =&\beta_1 + \sum_{i\geq 2}\gamma_i(\beta_i-\hat{\beta}_i) + \frac{X_1^T\epsilon}{|X_1|^2}.\end{aligned} \tag{9}$$

We only need to show the second term in (9) is of order $o_p(1/\sqrt{n})$. Bound the absolute value of that term with

$$\max_{i\geq 2}|\gamma_i|\cdot|\hat{\beta}_S-\beta_S|_1 \leq (c_1\lambda_n)\left(C_1s^*\lambda_n\right),$$

by assumption 1 and the l_1 control (5). That is then bounded by $O_p(s^*\lambda_n^2) = o_p(1/\sqrt{n})$ by assumption 3.

□

Remark 2. *With some careful manipulation the REC($3s^*, c_2$) condition as in assumption 2 can be reduced to REC(s^*, c_2). The proof would require an extra step of bounding $|\hat{\beta}_1-\beta_1|$ by $o_p(|\hat{\beta}_S-\beta_S|_1) + O_p(1/\sqrt{n})$.*

The ideas in the proofs for the two de-biasing estimators $\hat{\beta}_1^{(ZZ)}$ and $\hat{\beta}_1^{(one-step)}$ are similar. Ideally we want to run the regression

$$\arg\min_{b_1 \in \mathbb{R}} |Y - X_{-1}\beta_{-1} - b_1 X_1|^2. \tag{10}$$

That gives a perfectly efficient and unbiased estimator. However β_{-1} is not observed. It is natural to replace it with an estimator which is made globally close to the truth β_{-1} using penalized likelihood approach. As seen in the proof of Theorem 3, most of the work goes into establishing global l_1 control (5). The de-biasing estimator is then obtained by running an ordinary least squares regression like (10), replacing β_{-1} by some estimator satisfying (5), so that the solution to the least squares optimization is close to the solution of (10) with high probability.

2.2 Bayesian analogue of de-biasing estimators.

We would like to give a Bayesian analogue to the de-biasing estimators discussed above. As pointed out in the last section, it is essential to establish l_1 control on the vector $b_{-1} - \beta_{-1}$. Castillo et al. [3] and Gao et al. [5] have proposed priors that penalize sparsity of submodel dimension and provided theoretical guarantees such as LASSO-type contraction rates under the posterior distribution. This is the prior construction of Gao et al. [5, Section 3].

1. The size s of the dimension of the sub-model in the direction orthogonal to X_1 has probability mass function $\pi(s) \propto \exp(-Ds \log p)$.
2. $S|s \sim Unif(Z_s := \{S \subset \{1, ..., p\} : |S| = s, X_S \text{ is full rank}\})$.
3. Given the subset selection S, the coefficients b_S has density $f_S(b_S) \propto \exp(-\eta_n |X_S b_S|)$ for suitably chosen η_n.

Gao et al. [5] gave conditions under which we have a good l_1 posterior contraction rate.

Lemma 1. *(Corollary 5.4, [5]) If the design matrix X satisfies*

$$\kappa_0((2+\delta)s^*, X) = \inf_{|b|_0 \leq (2+\delta)s^*} \frac{\sqrt{s^*}|Xb|}{\sqrt{n}|b|_1} \geq c \tag{11}$$

for some positive constant c, δ, then there is constant $c_3 > 0$ and large enough $D > 0$ for which

$$\mathbb{P}_\beta \mu_Y \{|b - \beta|_1 > c_3 s^* \lambda_n\} \to 0,$$

We slightly modify the sparse prior of Gao et al. [5] to give good, asymptotically normal posterior behavior for a single coordinate. As we discussed in the last section, classical approaches to de-biasing exploit the idea of penalizing all coordinates except the one of interest. Our prior construction mimics that idea by putting the sparse prior only on b_{-1}.

2.3 The prior.

Denote the matrix projecting $\mathbb{R}^n$ to $span(X_1)$ by H. Under the model where $Y \sim \mathcal{N}(Xb, I_n)$, the likelihood function has the factorization

$$\begin{aligned}\mathcal{L}_n(b) =& \frac{1}{\sqrt{n}(2\pi)^{n/2}} \exp\left(-\frac{|Y - Xb|^2}{2}\right) \\ =& \frac{1}{\sqrt{n}(2\pi)^{n/2}} \exp\left(-\frac{|HY - HXb|^2}{2}\right) \exp\left(-\frac{|(I-H)Y - (I-H)Xb|^2}{2}\right).\end{aligned}$$

Write $W = (I - H)X_{-1}$ and reparametrize $b_1^* = b_1 + \sum_{i \geq 2} \gamma_i b_i$ with γ_i as defined in assumption 1. The likelihood $\mathcal{L}_n(b)$ can be rewritten as a constant multiple of

$$\exp\left(-\frac{|HY - b_1^* X_1|^2}{2}\right) \exp\left(-\frac{|(I-H)Y - Wb_{-1}|^2}{2}\right).$$

The likelihood factorizes into a function of b_1^* and b_{-1}. Therefore if we make b_1^* and b_{-1} independent under the prior, they will be independent under the posterior. We put a Gaussian prior on b_1^* to mimic the ordinary least square optimization step in the classical approaches. We put a sparse prior analogue to that of Gao et al. [5, section 3] on b_{-1}, using W as the design matrix in the prior construction. By lemma 1, b_{-1} is close to β_{-1} in l_1 norm with high posterior probability as long as $\kappa_o((2+\delta)s^*, W)$ is bounded away from 0.

We make b_1^* and b_{-1} independent under the prior distribution. The product distribution corresponds to a prior distribution on the original vector b. Note that under the prior distribution b_1 and b_{-1} are not necessarily independent.

This modified prior also has the effect of eliminating a bias term, in a fashion analogues to that of the two-step procedure $\hat{\beta}_1^{(ZZ)}$. The joint posterior distribution of b_1^* and b_{-1} factorizes into two marginals. In the X_1 direction, the posterior distribution of b_1^* is asymptotically Gaussian centered around $\frac{X_1^T Y}{|X_1|^2} = \beta_1^* + \frac{X_1^T \epsilon}{|X_1|^2}$. After we reverse the reparametrization we want the posterior distribution of b_1 to be asymptotically Gaussian centered around an efficient estimator $\hat{\beta}_1 = \beta_1 + \frac{X_1^T \epsilon}{|X_1|^2} + o_p(1/\sqrt{n})$. Therefore we need to show

$b_1^* - b_1$ is very close to $\beta_1^* - \beta_1$. That can be obtained from the l_1 control on $b_{-1} - \beta_{-1}$ under the posterior. In the next section we will give the proof to our main posterior asymptotic normality result (Theorem 2) in detail.

3 Proof of Theorem 2.

Since that prior and the likelihood of b_1^* are both Gaussian, we can work out the exact posterior distribution.

$$b_1^*|Y \sim \mathcal{N}\left(\frac{\sigma_n^2}{1+|X_1|^2\sigma_n^2}X_1^T Y, \frac{\sigma_n^2}{1+|X_1|^2\sigma_n^2}\right).$$

Since b_1^* and b_{-1} are independent under the posterior distribution, the above is also the distribution of b_1^* given Y and b_{-1}. That implies the distribution of $|X_1|(b_1 - \hat{\beta}_1)$ given Y and b_{-1} is

$$\mathcal{N}\left(|X_1|\left(\frac{\sigma_n^2}{1+|X_1|^2\sigma_n^2}X_1^T Y - \sum_{i\geq 2}\gamma_i b_i - \hat{\beta}_1\right), \frac{\sigma_n^2|X_1|^2}{1+|X_1|^2\sigma_n^2}\right). \tag{12}$$

Note that without conditioning on b_{-1}, the posterior distribution of b_1 is not necessarily Gaussian.

The goal is to show the bounded-Lipschitz metric between the posterior distribution of b_1 and $\mathcal{N}(\hat{\beta}_1, 1/|X_1|^2)$ goes to 0 under the truth. From Jensen's inequality and the definition of the bounded-Lipschitz norm we have

$$\begin{aligned}&\left\|\mathcal{L}(|X_1|(b_1-\hat{\beta}_1)|Y) - \mathcal{N}(0,1)\right\|_{BL}\\ \leq&\mu_Y^{b_{-1}}\left\|\mathcal{L}(|X_1|(b_1-\hat{\beta}_1)|Y, b_{-1}) - \mathcal{N}(0,1)\right\|_{BL}.\end{aligned}$$

For simplicity denote the posterior mean and variance in (12) as ν_n and τ_n^2 respectively. The bounded-Lipschitz distance between two normals $\mathcal{N}(\mu_1, \sigma_1^2)$ and $\mathcal{N}(\mu_2, \sigma_2^2)$ is bounded by $(|\mu_1 - \mu_2| + |\sigma_1 - \sigma_2|) \wedge 2$. Hence the above is bounded by

$$\mu_Y^{b_{-1}}\left(|\nu_n| \wedge 2\right) + \mu_Y^{b_{-1}}\left((|\tau_n - 1|) \wedge 2\right).$$

Therefore to obtain the desired convergence in (4), we only need to show

$$\mathbb{P}_\beta \mu_Y^{b_{-1}}\left(|\nu_n| \wedge 2\right) \to 0, \quad \text{and} \tag{13}$$

$$\mathbb{P}_\beta \mu_Y^{b_{-1}}\left((|\tau_n - 1|) \wedge 2\right) \to 0. \tag{14}$$

To show (13), notice that the integrand is bounded. Hence it is equivalent to show convergence in probability. Write

$$\begin{aligned}|\nu_n| =& \frac{\sigma_n^2|X_1|^d}{1+\sigma_n^2|X_1|^2}\left(\beta_1 + \frac{X_1^T\epsilon}{|X_1|^2} + \sum_{i\geq 2}\gamma_i\beta_i\right) \\ & - \sum_{i\geq 2}\gamma_i b_i - |X_1|\left(\beta_1 + \frac{X_1^T\epsilon}{|X_1|^2} + o_p\left(\frac{1}{\sqrt{n}}\right)\right) \\ \leq & \frac{|X_1|}{1+\sigma_n^2|X_1|^2}\left|\beta_1 + \frac{X_1^T\epsilon}{|X_1|} + \sum_{i\geq 2}\gamma_i\beta_i\right| + \sum_{i\geq 2}\gamma_i(\beta_i - b_i) + o_p(1). \quad (15)\end{aligned}$$

The first term is no longer random in b, and it can be made as small as we with now that it is decreasing in σ_n. If we set $\sigma_n^2 \gg |\beta|_1\lambda_n/|X_1|$, this term is of order $o_p(1)$.

For the second term, we will apply lemma 1 to deduce that this term also goes to 0 in $\mathbb{P}_\beta\mu_Y^{b_{-1}}$ probability. To apply the posterior contraction result we need to establish the compatibility assumption (11) on W.

Lemma 2. *Under assumption 1, 2, 3 and the constraint $|X_1| = O(\sqrt{n})$, the matrix $W = (I-H)X_{-1}$ satisfies*

$$\kappa_0((2+\delta)s^*, W) = \inf_{|b|_0\leq(2+\delta)s^*}\frac{\sqrt{s^*}|Wb|}{\sqrt{n}|b|_1} \geq c$$

for some $c, \delta > 0$.

We will prove the lemma after the proof of Theorem 2.

To show (14), Note that the integrand is not a random quantity. It suffices to show

$$|\tau_n - 1| = \left|\frac{\sigma_n^2|X_1|}{1+\sigma_n^2|X_1|^2} - \frac{1}{|X_1|}\right| \to 0.$$

That is certainly true for a $\{\sigma_n\}$ sequence chosen large enough. Combine (13), (14) and the bound on the bounded Lipschitz distance, we have shown

$$\mathbb{P}_\beta\left\|\mathcal{L}\left(|X_1|(b_1 - \hat{\beta}_1)|Y\right) - \mathcal{N}(0,1)\right\|_{BL} \to 0.$$

Proof of lemma 2. We will justify the compatibility assumption on W in two steps. First we will show that the compatibility assumption of the X

matrix follows from the REC assumption 2. Then we will show that the compatibility constant of X and W are not very far apart.

Let us first show that under assumption 2, there exist constants $0 < \delta < 1$ and $c > 0$, for which

$$\kappa_0((2+\delta)s^*, X) = \inf_{|b|_0 \leq (2+\delta)s^*} \frac{\sqrt{s^*}|Xb|}{\sqrt{n}|b|_1} \geq c.$$

Denote the support of g as S. We have

$$\begin{aligned}
\kappa_0((2+\delta)s^*, X) \geq & \inf_{|b|_0 \leq (2+\delta)s^*} \frac{1}{\sqrt{2+\delta}} \frac{|Xb|}{\sqrt{n}|b_S|} \\
\geq & \min_{\substack{J \subset [p], \\ |J| \leq 3s^*}} \inf_{\substack{b \neq 0, \\ |b_{J^C}|_1 \leq c_2 |b_J|_1}} \frac{|Xb|}{\sqrt{n}|b_J|} \\
= & \kappa(3s^*, c_2) > 0.
\end{aligned}$$

Now, under assumptions 1, 2 and 3, we will show that there exist constants $0 < \delta' < 1$ and $c' > 0$, for which

$$\kappa_0((2+\delta')s^*, W) \geq \kappa_0((2+\delta)s^*, X) + o(1).$$

For $g \in [R]^{p-1}$, we have

$$\begin{aligned}
|Wg| = & \left| X \begin{bmatrix} 0 \\ g \end{bmatrix} - \sum_{i \geq 2} \gamma_i g_i \right| \\
\geq & \left| X \begin{bmatrix} 0 \\ g \end{bmatrix} \right| - \lambda_n |g_1|
\end{aligned}$$

by assumption 1. Deduce that

$$\begin{aligned}
\kappa_0((2+\delta')s^*, W) = & \inf_{|b|_0 \leq (2+\delta)s^*} \frac{\sqrt{s^*}|Wb|}{\sqrt{n}|b|_1} \\
\geq & \kappa_0((2+\delta')s^* + 1, X) - \sqrt{\frac{s^*}{n}} \lambda_n \\
= & \kappa_0((2+\delta')s^* + 1, X) - \frac{\sqrt{s^* \log p}}{n}.
\end{aligned}$$

The second term if order $o(1)$ under assumption 3.

□

Acknowledgement

I would like to thank my advisor, Professor Pollard, for his expert advise and extraordinary support throughout this project.

References

[1] Peter J Bickel, Ya'acov Ritov, and Alexandre B Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *Annals of Statistics*, 37(4):1705–1732, 2009.

[2] Emmanuel Candes and Terence Tao. The Dantzig selector: Statistical estimation when p is much larger than n. *Annals of Statistics*, 35(35):2313–2351, 2007.

[3] Ismael Castillo, Johannes Schmidt-Hieber, Aad Van der Vaart, et al. Bayesian linear regression with sparse priors. *Annals of Statistics*, 43(5):1986–2018, 2015.

[4] R.M. Dudley. Speeds of metric probability convergence. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 22(4):323–332, 1972.

[5] Chao Gao, Aad W van der Vaart, and Harrison H Zhou. A general framework for Bayes structured linear models. *arXiv:1506.02174*, 2015.

[6] David Pollard. *A User's Guide to Measure Theoretic Probability*, volume 8 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, 2002.

[7] Robert Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.

[8] Aad W Van der Vaart. *Asymptotic Statistics*, volume 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge university press, 2000.

[9] Cun-Hui Zhang and Stephanie S Zhang. Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):217–242, 2014.